

TITLE OF THE INVENTION

**FAULT TOLERANT ATM-BASED DISTRIBUTED VIRTUAL TANDEM
SWITCHING SYSTEM AND METHOD**

INVENTORS

Haifeng BI
Tina M. SIGARTO

P18448.S01

FAULT TOLERANT ATM-BASED DISTRIBUTED VIRTUAL TANDEM
SWITCHING SYSTEM AND METHOD

CROSS-REFERENCE TO RELATED APPLICATIONS

5 This application is a continuation-in-part of pending U.S. Patent Application No. 09/287,092, filed April 7, 1999, to George C. ALLEN Jr. et al., entitled "ATM-based Distributed Virtual Tandem Switching System," which claims the benefit of U.S. Provisional Patent Application No. 60/083,640 filed on April 30, 1998, entitled "ATM-Based Distributed Virtual Tandem Switching System" to ALLEN et al., the disclosures of which are expressly incorporated herein by reference in their entireties.

10 BACKGROUND OF THE INVENTION

15 1. Field of the Invention

The present invention relates to the field of telecommunications. More particularly, the present invention relates to reliably constructing and operating asynchronous transfer mode (ATM)-based telecommunications networks.

20 2. Background Information

Tandem replacement using voice trunking over ATM (VTOA) technology, described in U.S. patent application no. 09/287,092, to George C. ALLEN Jr. et al., entitled "ATM-based Distributed Virtual Tandem Switching System," filed on April 7, 1999, is one application of an ATM distributed network system architecture. The architecture represents a new paradigm of networking that requires re-thinking of how to run networks. An important consideration is how to construct and operate the new ATM-based virtual tandem switch as reliably as possible and, definitely, no less reliably than current time division multiplexed (TDM) tandems.

25 The ATM-based virtual tandem system impacts system reliability. On one hand, the virtual tandem improves system reliability by distributing its components

geographically, localizing the impact of failures. On the other hand, a greater number of network elements is involved, and thus the number of occurrences of element failures may increase. Because the virtual tandem is designed to serve an entire metropolitan area, it is imperative for the virtual tandem's design to meet the highest level of survivability.

5 The present invention identifies potential failure points in the virtual tandem and provides solutions to reduce and survive failures. The solutions, in turn, place design and engineering requirements upon equipment vendors and companies operating the virtual tandem. It is therefore a primary object of the present invention to employ these requirements for use in the design of the network elements and for engineering the networks.

10 With reference to Fig. 1 of the drawings, standard call processing employs end offices 10 connected via tandem trunks 12, direct trunks 14, or both tandem trunks 12 and direct trunks 14. Each trunk 12, 14 is a digital service level 0 (DS0), operating at 64 kbps, that is transmitted between the switching offices 10 in a time division multiplexed manner. Each end office 10 connects to its neighboring end office 10 and the tandem office 16 using separate trunk groups. In this system, trunk groups are forecasted and pre-provisioned with dedicated bandwidth, which may lead to inefficiency and high operations cost.

15 20 A new voice trunking system using ATM technology has been proposed in U.S. patent application no. 09/287,092, entitled "ATM-Based Distributed Virtual Tandem Switching System." In this system, shown in Fig. 2, voice trunks from end office switches 20, 26 are converted to ATM cells by a first or second trunk inter-working function (T-IWF) device 22, 24. The T-IWFs 22, 24 are distributed to each end office 20, 26, and are controlled by a centralized control and signaling inter-

working function (CS-IWF) device 28. The CS-IWF 28 performs call control functions as well as conversion between the narrowband Signaling System No. 7 (SS7) protocol and a broadband signaling protocol. The T-IWFs 22, 24, CS-IWF 28, and the ATM network 30 form the ATM-based distributed virtual tandem switching system. According to this voice trunking over ATM (VTOA) architecture, trunks are no longer statistically provisioned as DS0 time slots. Instead, the trunks are realized through dynamically established switched virtual connections (SVCs), thus eliminating the need to provision separate trunk groups to different destinations, as done in TDM-based trunking networks.

10 BRIEF DESCRIPTION OF THE DRAWINGS

15 The present invention is further described in the detailed description that follows, by reference to the noted plurality of drawings by way of non-limiting examples of preferred embodiments of the present invention, in which like reference numerals represent similar parts throughout several views of the drawings, and in which:

20 Fig. 1 shows a conventional TDM telecommunications network architecture;

Fig. 2 shows a known virtual trunking over ATM telecommunications network architecture;

25 Fig. 3 shows a CS-IWF complex architecture, according to one aspect of the present invention;

Fig. 4 shows an end office architecture and its relationship with an ATM network, according to another aspect of the present invention; and

25 Fig. 5 shows an SMS connected to an ATM network, according to a further aspect of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

In view of the foregoing, the present invention is directed to improving the reliability of the VTOA system. The present invention identifies potential failure points of the VTOA system and provides various configurations to minimize the impact of failures.

According to an embodiment of the present invention, a control and signaling interworking function (CS-IWF) complex is provided for use within a VTOA system that communicates with mated signaling transfer points. The VTOA system includes an ATM network including multiple ATM switches and multiple trunk interworking functions (T-IWFs). The CS-IWF complex includes multiple CS-IWF units connected to at least two of the ATM switches and connected to at least one of the signaling transfer points. Each CS-IWF unit has multiple processors with at least one processor compensating for a failed processor. When a CS-IWF unit fails, at least one other CS-IWF unit compensates for the failed CS-IWF unit. The processor(s) compensating for the failed processor cooperate with the CS-IWF unit(s) compensating for the failed CS-IWF unit so that the CS-IWF complex survives simplex failures.

According to another embodiment, the CS-IWF complex includes multiple signaling link sets. Each link set connects the CS-IWF complex to one of the mated STPs. The CS-IWF complex may also include multiple signaling gateways that connect to each of the CS-IWF units. Each signaling gateway connects to one of the mated STPs. The multiple signaling gateways minimize isolation of the CS-IWF units when a link failure occurs. The CS-IWF complex also includes multiple ATM links that connect each CS-IWF unit to multiple ATM switches.

In one embodiment, the multiple processors within each CS-IWF unit operate

in an active/standby mode. Alternatively, the multiple processors operate in a load sharing mode.

Preferably, at least one of the CS-IWF units is located in a building separate from a building housing at least one other of the CS-IWF units. Furthermore, a single point code identifies the CS-IWF complex.

5 According to an embodiment of the present invention, an end office building is provided for use with a VTOA system. The VTOA system includes an ATM network including interconnected ATM switches, and at least one CS-IWF complex. The end office building includes multiple T-IWFs, which are part of the VTOA system. Each T-IWF has multiple processors with at least one processor compensating for a failed processor. Moreover, at least one T-IWF absorbs at least a portion of a failed T-IWF's workload. The end office building also includes a switch that distributes calls among the T-IWFs in a load sharing manner. Thus, the processor(s) compensating for the failed processor cooperate with the T-IWF(s) absorbing at least a portion of the failed T-IWF's workload so that the end office building survives simplex failures.

10 According to another embodiment, the end office building also includes at least one add/drop multiplexor (ADM) that connects the multiple T-IWFs to the ATM network. Preferably, each T-IWF also includes an optical interface for connecting to the ADM. The optical interface supports SONET 1+1 automatic protection switching.

15 In one embodiment, at least one T-IWF connects to a first ATM switch that is different from a second ATM switch to which another T-IWF connects. Thus, each end office building connects to multiple ATM switches so that if one ATM switch fails, the end office building remains connected to the ATM network.

According to another embodiment of the present invention, a method is provided for recovery from a failing ATM link in a VTOA system. The VTOA system includes an ATM network having multiple ATM switches interconnected by ATM links, multiple T-IWFs, and at least one CS-IWF complex. The method includes delaying recovery action in the ATM network for a predetermined duration while SONET recovery of the link is attempted. If the SONET recovery is successful, a call path through the ATM network stays up. If the SONET recovery is unsuccessful, existing calls carried by the failed ATM link are dropped. Preferably, the predetermined duration is 100 milliseconds.

According to a further embodiment of the present invention, a switch management system (SMS) is provided for use within a VTOA system. The VTOA system includes an ATM network including a least one ATM switch, multiple T-IWFs, and at least one CS-IWF complex. The switch management system includes multiple switch management system units. At least one of the switch management system units is a backup unit for at least one primary switch management system unit. Each switch management system unit provides application redundancy within itself. Consequently, the switch management system survives simplex failures.

Preferably, the primary switch management system unit is located in a building separate from a building housing the backup switch management system unit. In addition, each switch management system unit is connected to multiple ATM switches.

According to yet another embodiment of the present invention, a method is provided for restoring functions of a failed switch management system operating within a VTOA system. The VTOA system includes an ATM network including multiple ATM switches, multiple T-IWFs, and at least one CS-IWF complex. The

method includes restoring essential surveillance of the VTOA system; restoring billing functions of the VTOA system; restoring full surveillance capability of the VTOA system; restoring configuration management of the VTOA system; and restoring performance management of the VTOA system.

5

10

15

20

25

According to yet another embodiment of the present invention, a VTOA system includes an ATM network having multiple interconnected ATM switches. Also provided are multiple mated signaling transfer points that communicate with the VTOA system. The VTOA system also includes at least one CS-IWF complex including multiple CS-IWF units connected to at least two of the ATM switches and connected to at least one of the signaling transfer points. Each unit has multiple processors that share load resulting from failure of one of the processors. In addition, at least one end office building is provided for interaction with the VTOA system. Each end office building includes multiple T-IWFs, each having multiple processors. A switch is also provided in the end office building(s) to distribute calls among the multiple T-IWFs in a load sharing manner. The VTOA system also includes a switch management system including multiple switch management system units. At least one of the switch management system units is a backup unit for at least one primary switch management system unit. Each switch management system unit provides application redundancy within itself. Consequently, the VTOA system survives simplex failures. According to another embodiment, there are at least two completely disjointed routes between any two end points.

According to yet another embodiment of the present invention, a method is provided for communicating over a VTOA system. The VTOA system includes a CS-IWF complex, an ATM network containing multiple ATM switches, and a signaling network. Multiple end office buildings are provided for interaction with the

VTOA system. Each end office building includes a switch and multiple T-IWFs, which are part of the VTOA system. The method includes transmitting a signal from the switch to the T-IWFs in a load sharing manner; and transmitting from the T-IWFs to the ATM switches in a load sharing manner. Thus, communications survive a simplex failure in the VTOA system.

5 According to the present invention, the following failure points are analyzed: CS-IWF failure; T-IWF failure; ATM network failure; and SMS failure. Each of these failure scenarios is discussed below along with survivability measures to protect against and survive these failures. Each solution, discussed below, guarantees that the network will survive all simplex failures. A simplex failure is the occurrence of a single network element failure, in contrast to simultaneous failures of multiple network elements. To survive means that, in the event of a simplex failure, the network must be able to continue to operate and recover on its own either to its normal or to a compromised level of performance.

10

CS-IWF

15 The CS-IWF is the most critical element of the new virtual tandem because failure of the CS-IWF impacts the entire serving area. Thus, failure of the CS-IWF is not acceptable, and therefore the CS-IWF requires the highest level of reliability. Exemplary CS-IWF units include the Connection Gateway from Lucent Technologies Inc, and the Succession Call Server, from Nortel Networks Corporation.

20

25 Figure 3 shows a design for a reliable CS-IWF complex 300 that includes multiple CS-IWF units 310, 320, 330. According to the present invention, each CS-IWF complex 300 serving a metropolitan area is assigned a single point code, regardless of how many individual CS-IWF units 310, 320, 330 the complex 300

P18448.S01

contains. In Figure 3, a general case is depicted where the CS-IWF complex 300 includes of n CS-IWFs 310, 320, 330 for reasons of reliability or processing capacity. A special case occurs when $n = 2$, in which case two CS-IWFs 310, 320 operate in a load sharing or active/standby mode.

According to an object of the present invention, each CS-IWF unit 310, 320, 330 must be highly reliable. To achieve this objective, redundant processors are provided within each CS-IWF 310, 320, 330 for protection against processor failure. In Figure 3, processor 0 311, 321, 331 and processor 1 312, 322, 332 are shown, although one of ordinary skill in the art will realize that more processors can be added without departing from the scope of the present invention. The redundant processors may operate in an active/standby mode or in a load sharing mode.

Each CS-IWF complex 300 must contain spare capacity for protection. The specific architecture of the CS-IWF complex 300 dictates the spare processing capacity required. For example, in a complex 300 where $n = 2$, if one CS-IWF 310 fails, the remaining CS-IWF 320 must be able to handle the load of the CS-IWF 310 that failed. If three CS-IWFs 310, 320, 330 are provided, any two remaining CS-IWFs 320, 330 should be able to handle the load of the failed CS-IWF 310. Thus, a CS-IWF complex 300 must contain at least two CS-IWF units 310, 320. In general, in a CS-IWF complex 300 of n units, up to k ($k \geq 1$) out of the n CS-IWF units 310, 320, 330 must be provided for the purpose of protection. The objective is that the loss of one CS-IWF 310 unit has no impact on the call handling capacity of the CS-IWF complex 300 as a whole. In the active/standby mode, $n - k$ CS-IWFs 310, 320 are active, and k operate in standby mode. In the load-sharing mode, all n CS-IWFs 310, 320, 330 run at levels less than maximum such that if one of the CS-IWFs 310 should fail, its processing load can be absorbed by the remaining CS-IWFs 320, 330.

25

In an embodiment, all components of a CS-IWF complex 300 are not provided at the same physical location. As a result, the loss of one physical location does not shut down the entire network. In this embodiment, the components can be connected via either dedicated or networked links. All components of the CS-IWF complex 300 are NEBS level 3 compliant. See GR-63-CORE, (*Network Equipment – Building System [NEBS] Requirements: Physical Protection*)—A module of LSSGR, FR-64; TSGR, FR-440; and NEBSFR, FR-2065; GR-1089-CORE, (*Electromagnetic Compatibility and Electrical Safety Generic Criteria for Network Telecommunications Equipment*)—A module of LSSGR, FR-64; and TSGR, FR-440; and SR-NWT-002550, (*Technical Considerations for NEBS-2000*), the disclosures of which are expressly incorporated by reference in their entireties, for more about NEBS level 3.

As is well known, STPs typically include mated pairs 370, 380 operating in a load-sharing manner to increase reliability. To take advantage of the mated STP's reliability and to further improve the CS-IWF complex's 300 reliability, a CS-IWF complex 300 maintains, at a minimum, two signaling link sets, one with each of the local mated STPs 370, 380. Moreover, within a CS-IWF complex 300, the CS-IWF units 310, 320, 330 are connected with the signaling link sets in a configuration that minimizes the possibility of any CS-IWF unit 310, 320, 330 being isolated from the signaling network as a result of a link or link-set failure. A non-limiting example of such interconnection using signaling gateways 350, 360 is shown in Figure 3. The CS-IWF complex must support low-speed signaling links (e.g., 56 kbps to operate with SS7) with the ability to migrate to high-speed links (e.g., T1 or T1 ATM).

Each CS-IWF complex 300 bridges between narrowband and broadband signaling. For example, the narrowband signaling may be in the form of SS7 ISUP

messages, and the broadband signaling may be standard-based broadband signaling, for example ATM UNI or PNNI.

The signaling gateways 350, 360 are part of the CS-IWF complex 300 and distribute SS7 signaling to each CS-IWF unit 310, 320, 330. The signaling gateways 5 350, 360 form an interconnection network that connects STPs 370, 380 to CS-IWF units 310, 320, 330. In other words, the signaling gateways are a distribution vehicle. An exemplary signaling gateway is the Connection Gateway Signaling Node, manufactured by Lucent Technologies, Inc.

In another embodiment, each CS-IWF 310, 320, 330 maintains an ATM link 10 with two different ATM switches 390, 395 in the ATM network so that the CS-IWF complex 300 can communicate with T-IWFs, other CS-IWFs and the SMS (not shown in Fig. 3). Preferably the ATM switches 390, 395 are on separate SONET rings. Further, it is not necessary for all CS-IWFs 310, 320, 330 to connect to the same two ATM switches 390, 395. According to this embodiment, the well known 15 1 + 1 automatic protection switching (APS) is not required.

T-IWF

It is an object of the present invention that each T-IWF is highly reliable, i.e., 20 nearly always available. Therefore, according to one embodiment, the critical components (e.g., processor) within a T-IWF are redundant to achieve this object. Exemplary T-IWFs include the 7R/E Trunk Access Gateway, from Lucent Technologies Inc.; and the Succession Multi-service Gateway 4000 (MG 4000), from Nortel Networks Corporation.

Figure 4 illustrates an exemplary architecture of an end office building 400 and 25 its relationship to an ATM network 475. In Figure 4, a class 5 end office building

(EO) 400 includes a switch 402, associated T-IWFs 404, 406, 408, which should be NEBS level 3 compliant, and a SONET add/drop multiplexer (ADM) 410. Exemplary switches include class 5 switches such as: the Lucent Technologies Inc. 1AESS; the Lucent Technologies Inc. 5ESS; the Ericsson AXE-10; and/or the Northern Telecom (Nortel) DMS-100 switches. In Figure 4 a general case is shown where multiple T-IWF units 404, 406, 408, are deployed in an end office building 400 for reasons of reliability or capacity. Although the switch 402 and ATM switches 420, 430 shown in Figure 4 are not co-located in the same end office building 400, such co-location may occur in other end office buildings.

According to the present invention, a class 5 switch having traffic volume requiring only one T-IWF is still connected with two T-IWFs for protection. Consequently, loss of one T-IWF does not isolate the class 5 switch. Furthermore, a class 5 switch must be able to maintain as few as one trunk group regardless of the number of T-IWFs by which it is served. According to an embodiment, the class 5 switch 402 distributes its calls among the T-IWFs 404, 406, 408 in a load-sharing manner. Thus, loss of one of the T-IWFs 404, 406, 408 may degrade the trunking capacity of the class 5 switch 402, but will not isolate the class 5 switch 402.

The optical interface on the T-IWF 404, 406, 408 for connecting with the SONET add/drop multiplexer (ADM) 410 (or an ATM switch 420, 430 when an ATM switch 420, 430 is located in the same end office building 400) supports the SONET 1+1 Automatic Protection Switching (APS) scheme, although deployment of this feature is optional. The ADM 410 connects the T-IWFs 404, 406, 408 to the ATM switches 420, 430 at the ATM layer via links 440, 450. The links 440, 450 are preferably SONET rings that connect the T-IWFs 404, 406, 408 to ATM switches 420, 430 in a load sharing manner. Exemplary ADMs are manufactured by Fujitsu, 420, 430 in a load sharing manner. Exemplary ADMs are manufactured by Fujitsu,

Lucent Technologies Inc., and Nortel.

Each T-IWF 404, 406, 408 serving a given end office building 400 does not connect to the ATM network at the same ATM switch 420, 430. Rather, each T-IWF 404, 406, 408 is single-homed to an ATM switch 420, 430, while each end office building 400 is multi-homed to multiple ATM switches 420, 430, preferably on separate SONET rings 440, 450. This configuration prevents the end office building 400 from becoming totally disconnected from the ATM network in the event of an ATM host 420, 430 failure, although the T-IWFs 404, 406, 408 connected to the failed ATM switch 420, 430 may be impacted. Preferably, the SONET transport network 440, 450 employs uni-directional path switched ring (UPSR) or bi-directional line switched ring (BLSR). In order to protect against both ATM and SONET layer failures, the ATM layer virtual path protection capability in a SONET/ATM hybrid transport node may be supported. Further, the ATM VP ring functionality may be integrated into the T-IWF 404, 406, 408 and the ATM switches 420, 430. According to another embodiment, ATM virtual path (VP) ring functionality is either integrated with or separated from the T-IWF 404, 406, 408 to achieve the potential benefits of ATM layer VP protection and transport layer efficiency.

ATM network

An example ATM network environment, as relevant to the analysis of the failure scenarios, is now discussed. However, if an ATM network having a different configuration is provided, alternate CS-IWF, T-IWF, SMS, etc. configurations from those presently described may be preferred. In the exemplary ATM network, ATM switches are not protected with redundant ATM switches; nor are ATM virtual connections carrying bearer channels protected with redundant virtual connections.

The 1 + 1 protection of user network interface (UNI) interfaces on ATM switches is not universally deployed. In the absence of the ATM VP ring capability, no attempt (such as virtual connection re-routing) will be made at the ATM layer to save calls in progress that are impacted by an ATM equipment or link failure.

5

The 1 + 1 protection enables SONET recovery by directing traffic from a failed ring, e.g., a cut ring, to a properly functioning ring. That is, two devices are connected by two rings (one active ring and one standby ring) in the usual manner. When the active ring fails, the standby ring is activated.

10

In order to eliminate single points of total failure in the ATM network, the network must be constructed so that between any two end points at least two completely disjointed routes traverse the ATM network. Consequently, the ATM switches are able to intelligently route calls in this network, e.g., by balancing the call load between disjointed routes to reduce the impact of failure as well as to improve network performance. The balanced intelligent routing is performed in a known manner.

15

An ATM link failure occurs as result of a transport facility failure, such as a fiber cut. The ATM network therefore relies on known protection schemes in the transport network to recover from such a failure. In the event that the ATM switch detects a link failure, the ATM switch delays recovery action for a predetermined time period, preferably 100 ms, during which time the SONET layer recovery is attempted.

20

If the transport layer successfully recovers, then the impact of the failure will only be a momentary degradation of the voice connections, and the connected call paths stay up. If the transport layer fails to recover from such a link failure, then existing calls being carried by the link are dropped. The ATM switches on both ends of the failed link will then flag the associated ports as unavailable, and future calls will

25

automatically avoid the failed link until it is repaired. After the link is repaired, no manual intervention is required in order for traffic to resume using that link, as is well known.

According to the present invention, ATM switching equipment failures only include failures of non-redundant components, such as un-protected interface cards, or a whole switch. Exemplary ATM switches include the MainStreetXpress 36170 Multiservices Switch or 670 RSP, both manufactured by Newbridge Networks Corporation; the GX 550 Smart Core ATM Switch, manufactured by Lucent Technologies Inc.; and the Passport 15000 Multiservice Switch, manufactured by Nortel Networks Corporation. In an embodiment, common equipment in an ATM switch, such as the switching fabric, the control processor, the power supply, wiring, fuses, alarms, etc. are redundant. Consequently, failure of one such unit has no impact on the operation or the performance of the ATM switch.

Redundant interface cards, however, are not provided. Thus, when an un-protected interface card or port fails, calls being carried by that interface card or port are dropped. The ATM switches connected via the failed interface card or port will then flag the associated card or port as unavailable, and future calls will automatically avoid using the failed card or port until it is repaired. After the card or port is repaired, minimal manual intervention will be required in order for traffic to resume using that card or port, as is well known.

In the exemplary ATM network, redundant ATM switches are not provided. In other words, an ATM switch does not have a standby. Thus, in the event of a total ATM switch failure, such as loss of the building, calls being carried by the ATM switch are dropped. The ATM network will then flag the failed switch as unavailable, and future calls will automatically avoid this switch until it is repaired. After the ATM

switch is repaired, minimal manual intervention will be required in order for traffic to resume using that ATM switch, as is well known.

SMS

5 Figure 5 shows a single switch management system (SMS) unit 500. The SMS is the element layer manager of the ATM-based virtual tandem. It communicates with the T-IWFs and the CS-IWF, and the legacy operation support systems (OSS). Essentially, it controls management of the distributed switch and acts as a man-
10 machine interface enabling a human user to view and control the overall behavior of the VTOA. According to one embodiment, it communicates with other network management systems involved in the virtual tandem, such as the operation support system of the ATM network. The SMS can be located either in a central office or in a data center, and should be NEBS level 3 compliant. Exemplary SMSs include the
15 OneLink Manager, from Lucent Technologies Inc., and the Succession Network Manager, from Nortel Networks Corporation.

15 According to an embodiment, the SMS includes a primary SMS unit 500 and a backup SMS unit (not shown) that takes over if the primary SMS unit 500 fails. That is, the primary and the backup SMS units operate in an active/standby mode. The
20 backup SMS unit may support multiple primary SMS units 500, as dictated by engineering and operational network requirements, and must be located in a different building from the building housing the primary SMS unit 500.

25 As seen in Figure 5, each SMS unit 500 maintains dual ATM links 510, 520 with two different ATM switches 530, 540, preferably on separate SONET rings. The dual links 510, 520 allow control communications with the backup SMS unit, the T-IWFs, and the CS-IWF. In other words, each switch management system unit has

P18448.S01

management connectivity to other VTOA system elements provided by paths through the ATM network.

5 Each SMS unit must provide application redundancy within itself, with automatic, transparent switch over in the event of failure of the active SMS application. Redundancy may be accomplished by providing a backup processor in each physical platform and/or providing backup software applications. For example, two applications may run side by side in a single processor, or separately in two processors, or the second application may begin in the second processor when the first application fails. Consequently, if part of one physical platform fails, the remaining portion of the physical platform is configured so that it can compensate for the failed portion.

10 In the event of failure of the active SMS unit 500, the switch of its load to the other unit must be accomplished with minimal manual actions, ideally no actions and preferably no more than a system administrator approaching the physical platform and issuing necessary instructions to the SMS through a computer terminal. Failure of the active unit must have minimal impact on an operations user of an SMS graphical user interface (GUI). For example, the operations user should not have to re-boot the computer or re-log in to the computer to continue using the GUI. Slower processing of commands is acceptable, and alarms and/or notifications of the switchover are necessary.

15 20 If the active SMS unit 500 fails, the SMS as a system restores its operation in the following sequence:

25

1. Essential surveillance such as status and critical alarms;
2. Billing functions;
3. Full surveillance capability;

4. Configuration management;
5. Performance management.

Essential surveillance refers to capabilities such as determining whether the VTOA switch is functional or non-functional, and determining the overall health of individual components of the network. Billing is self explanatory. Full surveillance refers to capabilities such as viewing all state changes within the system, viewing alarms and events, e.g., a card within a component that failed, etc. Configuration management refers to capabilities such as rearranging equipment and adding new connections. Performance management refers to capabilities such as collecting data for such tasks as engineering or growth of the network.

10 Each SMS unit has its own continually updated database. Each database is synchronized with the other VTOA databases. The database enables the five functions discussed above and includes such information as the system users, networking software, the network inventory, security, etc. Awareness of the network topology is also provided by the database.

15 Although the invention has been described with reference to several exemplary embodiments, it is understood that the words that have been used are words of description and illustration, rather than words of limitation. Changes may be made within the purview of the appended claims, as presently stated and as amended, without departing from the scope and spirit of the invention in its aspects. Although 20 the invention has been described with reference to particular means, materials and embodiments, the invention is not intended to be limited to the particulars disclosed; rather, the invention extends to all functionally equivalent structures, methods, and uses such as are within the scope of the appended claims.

25 In accordance with various embodiments of the present invention, the methods

described herein are intended for operation as software programs running on a computer processor. Dedicated hardware implementations including, but not limited to, application specific integrated circuits, programmable logic arrays and other hardware devices can likewise be constructed to implement the methods described herein. It should also be noted that the software implementations of the present invention can be stored on a tangible storage medium such as a magnetic or optical disk, read-only memory or random access memory and be produced as an article of manufacture.